

Dialog state tracking, a machine reading approach using Memory Network

Julien Perez

Xerox Research Centre Europe
Grenoble, France

julien.perez@xrce.xerox.com

Fei Liu *

The University of Melbourne
Victoria, Australia

fliu3@student.unimelb.edu.au

Abstract

In an end-to-end dialog system, the aim of dialog state tracking is to accurately estimate a compact representation of the current dialog status from a sequence of noisy observations produced by the speech recognition and the natural language understanding modules. This paper introduces a novel method of dialog state tracking based on the general paradigm of machine reading and proposes to solve it using an End-to-End Memory Network, MemN2N, a memory-enhanced neural network architecture. We evaluate the proposed approach on the second Dialog State Tracking Challenge (DSTC-2) dataset. The corpus has been converted for the occasion in order to frame the hidden state variable inference as a question-answering task based on a sequence of utterances extracted from a dialog. We show that the proposed tracker gives encouraging results. Then, we propose to extend the DSTC-2 dataset and the definition of this dialog state task with specific reasoning capabilities like counting, list maintenance, yes-no question answering and indefinite knowledge management. Finally, we present encouraging results using our proposed MemN2N based tracking model.

1 Introduction

One of the core components of state-of-the-art and industrially deployed dialog systems is a dialog state tracker. Its purpose is to provide a compact representation of a dialog produced from past user inputs and system outputs which is called the dialog state. The dialog state summarizes the infor-

mation needed to successfully maintain and finish a dialog, such as users' goals or requests. In the simplest case of a so-called *slot-filling schema*, the state is composed of a predefined set of variables with a predefined domain of expression for each of them. As a matter of fact, in the recent context of end-to-end trainable machine learnt dialog systems, state tracking remains a central element of such architectures (Wen et al., 2016). Current models, mainly based on the principle of discriminative learning, tend to share three common limitations. First, the tracking task is performed using a fixed window of the past dialog utterances as support for decision. Second, the possible correlations between the set of tracked variables are not leveraged and individual trackers tend to be learnt independently. Third, the tracking task is summarized as the capability of inferring values for a predefined set of latent variables. Starting from these observations, we propose to formalize the task of state tracking as a particular instance of machine reading problem. Indeed, these formalization and the proposed resolution model called MemN2N (Weston et al., 2015) allow to define a tracker that is able to decide at the utterance level on the basis of the current entire dialog. Indeed, the model learns to focus its attention on the meaningful parts of the dialog regarding the currently asked slot and can eventually capture possible correlation between slots. As far as our knowledge goes, it is the first attempt to explicitly frame the task of dialog state tracking as a machine reading problem. Finally, such formalization allows for the implementation of approximate reasoning capability that has been shown to be crucial for any machine reading tasks (Weston et al., 2015) while extending the task from slot instantiation to question answering. This paper is structured as follows, Section 2 recalls the main definitions associated to transactional dialogs and describes

*Work carried out as an intern at XRCE

the associated problem of statistical dialog state tracking with both the generative and discriminative approaches. At the end of this section, the limitations of the current models in terms of necessary annotations and reasoning capabilities are addressed. Then, Section 3 depicts the proposed machine reading model for dialog state tracking and proposes to extend a state of the art dialog state tracking dataset, *DSTC-2*, to several simple reasoning capabilities. Section 4 illustrates the approach with experimental results obtained using a state of the art benchmark for dialog state tracking.

2 Dialog state tracking

2.1 Main Definitions

A dialog state tracking task is formalized as follows: at each turn of a dyadic dialog, the dialog agent chooses a dialog act d to express and the user answers with an utterance u . In the simplest case, the dialog state at each turn is defined as a distribution over a set of predefined variables, which define the structure of the state (Williams et al., 2005). This classic state structure is commonly called *slot filling* or *semantic frame*. In this context, the state tracking task consists of estimating the value of a set of predefined variables in order to perform a procedure or transaction which is the purpose of the dialog. Typically, a natural language understanding module processes the user utterance and generates an N-best list $o = \{(d_1, f_1), \dots, (d_n, f_n)\}$, where d_i is the hypothesized user dialog act and f_i is its confidence score. Various approaches have been proposed to define dialog state trackers. The traditional methods used in most commercial implementations use hand-crafted rules that typically rely on the most likely result from an NLU module (Yeh et al., 2014) and hardly models uncertainty. However, these rule-based systems are prone to frequent errors as the most likely result is not always the correct one (Williams, 2014).

More recent methods employ statistical approaches to estimate the posterior distribution over the dialog states allowing them to leverage the uncertainty of the results of the NLU module. In the simplest case where no ASR and NLU modules are employed, as in a text based dialog system (Henderson et al., 2013), the utterance is taken as the observation using a so-called bag of words representation. If an NLU module is available, stan-

dardized dialog act schemes can be considered as observations (Bunt et al., 2010). Furthermore, if prosodic information is available from the ASR component of the dialog system (Milone and Rubio, 2003), it can also be considered as part of the observation definition. A statistical dialog state tracker maintains, at each discrete time step t , the probability distribution over states, $b(s_t)$, which is the system’s *belief* over the state. The actual slot filling process is composed of the cyclic tasks of *information gathering* and integration, in other words – *dialog state tracking*. In such framework, the purpose is to estimate as early as possible in the course of a given dialog the correct instantiation of each variable. In the following, we will assume the state is represented as a set of variables with a set of known possible values associated to each of them. Furthermore, in the context of this paper, only the bag of words has been considered as an observation at a given turn but dialog acts or detected named entity provided by an SLU module could have also been incorporated.

Two statistical approaches have been considered for maintaining the distribution over a state given sequential NLU output. First, the discriminative approach aims to model the posterior probability distribution of the state at time $t + 1$ with regard to state at time t and observations $z_{1:t}$. Second, the generative approach attempts to model the transition probability and the observation probability in order to exploit possible interdependencies between hidden variables that comprise the dialog state.

2.2 Generative Dialog State Tracking

A generative approach to dialog state tracking computes the belief over the state using Bayes’ rule, using the belief from the last turn $b(s_{t-1})$ as a prior and the likelihood given the user utterance hypotheses $p(z_t|s_t)$, with z_t the observation gathered at time t . In prior works (Williams et al., 2005), the likelihood is factored and some independence assumptions are made: $b_t \propto \sum_{s_{t-1}, z_t} p(s_t|z_t, s_{t-1})p(z_t|s_{t-1})b(s_{t-1})$. A typical generative model uses a factorial hidden Markov model (Ghahramani and Jordan, 1997). In this family of approaches, scalability is considered as one of the main issues. One way to reduce the amount of computation is to group the states into partitions, as proposed in the Hidden Information State (HIS) model (Gasic and Young, 2011). Other

approaches to cope with the scalability problem in dialog state tracking is to adopt a factored dynamic Bayesian network by making conditional independence assumptions among dialog state components, and then using approximate inference algorithms such as loopy belief propagation (Thomson and Young, 2010) or a blocked Gibbs sampling as (Raux and Ma, 2011). To cope with such limitations, discriminative methods of state tracking presented in the next part of this section aim at directly model the posterior distribution of the tracked state using a chosen parametric form.

2.3 Discriminative Dialog State Tracking

The discriminative approach of dialog state tracking computes the belief over a state via a parametric model that directly represents the belief $b(s_{t+1}) = p(s_{t+1}|s_t, z_t)$. For example, Maximum Entropy has been widely used in the discriminative approach (Metallinou et al., 2013). It formulates the belief as follows: $b(s) = P(s|x) = \eta \cdot e^{w^T \phi(x,s)}$, where η is the normalizing constant, $x = (d_1^u, d_1^m, s_1, \dots, d_t^u, d_t^m, s_t)$ is the history of user dialog acts, $d_i^u, i \in \{1, \dots, t\}$, the system dialog acts, $d_i^m, i \in \{1, \dots, t\}$, and the sequence of states leading to the current dialog turn at time t . Then, $\phi(\cdot)$ is a vector of feature functions on x and s . Finally, w is the set of model parameters to be learned from annotated dialog data. Finally, deep neural models, performing on a sliding window of features extracted from previous user turns, have also been proposed in (Henderson et al., 2014c; Mrksic et al., 2016). Of the current literature, this family of approaches have proven to be the most efficient for publicly available state tracking datasets. Recently, deep learning based models implementing this strategy (Mrksic et al., 2016; Henderson et al., 2014a; Williams et al., 2016) have shown state of the art results. This approach tends to leverage unsupervised training word representation (Mikolov et al., 2013).

2.4 Current Limitations

Using error analysis (Henderson et al., 2014b), three limitations can be observed in the application of these inference approaches. First, current models tend to fail at considering long-tail dependencies that occurs on dialogs. For example, coreferences, inter-utterances informations and correlations between slots have been shown to be difficult to handle even with the usage of recurrent network models (Henderson et al., 2014a). To illustrate the




Figure 1: T-SNE transformation of the final state of DSTC-2 train set.

inter-slot correlation, Figure 1 depicted the t-SNE (van der Maaten and Hinton, 2008) projected final state of the dialog of the DSTC-2 training set. On the other hand, reasoning capabilities, as required in machine reading applications (Poon and Domingos, 2010; Etzioni et al., 2007; Berant et al., 2014; Weston et al., 2015) remain absent in these classic formalizations of dialog state tracking. Finally, tracking definition is limited to the capability to instantiate a predefined set of slots. In the next section, we present a model of dialog state tracking that aims at leveraging the current advances of MemN2N, a memory-enhanced neural networks and their approximate reasoning capabilities that seems particularly adapted to the sequential, long range dependency equipped and sparse nature of complex dialog state tracking tasks. Furthermore, this model allows to relax the hypothesis of strict utterance-level annotation that does not corresponds to common practices in industrial applications of transactional conversational user interfaces where annotations tend to be placed at a multi-utterance level or full-dialog level only.

3 Machine Reading Formulation of Dialog State Tracking

We propose to formalize the dialog state tracking task as a machine reading problem (Etzioni et al., 2007; Berant et al., 2014). In this section, we recall the main definitions of the task of machine reading, then describes the MemN2N, a memory-enhanced neural network architectures proposed to handle such tasks in the context of dialogs. Finally, we formalize the task of dialog state tracking as a machine reading problem and propose to

solve it using a memory-enhanced neural architecture of inference.

3.1 Machine Reading

The task of textual understanding has recently been formulated as a supervised learning problem (Kumar et al., 2015; Hermann et al., 2015). This task consists in estimating the conditional probability $p(a|d, q)$ of an answer a to a question q where d is a document. Such an approach requires a large training corpus of {Document - Query - Answer} triples and until now such corpora have been limited to hundreds of examples (Richardson et al., 2013). In the context of dialog state tracking, it can be understood as the capability of inferring a set of latent values l associated with a set of variables v related to a given dyadic or multi-party conversation d , from direct correlation and/or reasoning, using the course of exchanges of utterances, $p(l|d, v)$.

State updates at an utterance-level are rarely provided off-the-shelf from a production environment. In these environments, annotation is often performed afterward for the purpose of logging, monitoring or quality assessment. In the limit cases, as in human-to-human dialog systems, dialog-level annotations remains a common practice of annotation especially in personal assistance, customer care dialogs and, in a more general sense, industrial application of transactional conversational user interfaces. Another frequent setting consist of informing the state after a given number of utterance exchange between the locutors. So an additional effort of specific annotation is often needed in order to train a state of the art statistical state tracking model (Henderson et al., 2014b). In that sense, formalizing dialog state tracking at a sub-dialog level in order to infer hidden state variables with respect to a list of utterances started from the first one to any given utterance of a given dialog seems particularly appropriate. In the context of dialog state tracking challenges, the *DSTC-4* dialog corpus have been designed in such purpose but only consists of 22 dialogs. Concerning the *DSTC-2* corpus, the training data contains 2207 dialogs (15611 turns) and the test set consists of 1117 dialogs (Williams et al., 2016). This dataset is more suitable for our experiments.

For these reasons, the machine reading paradigm becomes a promising formulation for

the general problem of dialog state tracking. Furthermore, current approaches and available datasets for state tracking do not explicitly cover reasoning capabilities such as temporal and spatial reasoning, counting, sorting and deduction. We suggest that in the future dataset dialogs expressing such specific abilities should be developed. In this last part, several reasoning enhancements are suggested to the *DSTC-2* dataset.

3.2 End-to-End Memory Networks

The MemN2N architecture, introduced by (Weston et al., 2015), consists of two main components: supporting memories and final answer prediction. Supporting memories are in turn comprised of a set of input and output memory representations with memory cells. The input and output memory cells, denoted by m_i and c_i , are obtained by transforming the input context x_1, \dots, x_n (i.e a set of utterances) using two embedding matrices A and C (both of size $d \times |V|$ where d is the embedding size and $|V|$ the vocabulary size) such that $m_i = A\Phi(x_i)$ and $c_i = C\Phi(x_i)$ where $\Phi(\cdot)$ is a function that maps the input into a bag of dimension $|V|$.

Similarly, the question q is encoded using another embedding matrix $B \in \mathbb{R}^{d \times |V|}$, resulting in a question embedding $u = B\Phi(q)$. The input memories $\{m_i\}$, together with the embedding of the question u , are utilized to determine the relevance of each of the stories in the context, yielding in a vector of attention weights

$$p_i = \text{softmax}(u^\top m_i) \quad (1)$$

where $\text{softmax}(a_i) = \frac{e^{a_i}}{\sum_i e^{a_i}}$. Subsequently, the response o from the output memory is constructed by the weighted sum:

$$o = \sum_i p_i c_i \quad (2)$$

Other models of parametric encoding for the question and the document have been proposed in (Kumar et al., 2015). For the purpose of this presentation, we will keep with definition of Φ .

For more difficult tasks requiring multiple supporting memories, the model can be extended to include more than one set of input/output memories by stacking a number of memory layers. In this setting, each memory layer is named a hop and the $(k+1)^{\text{th}}$ hop takes as input the output of the k^{th} hop:

$$u^{k+1} = o^k + u^k \quad (3)$$

Lastly, the final step, the prediction of the answer to the question q , is performed by

$$\hat{a} = \text{softmax}(W(o^K + u^K)) \quad (4)$$

where \hat{a} is the predicted answer distribution, $W \in \mathbb{R}^{|V| \times d}$ is a parameter matrix for the model to learn and K the total number of hops.

Two weight tying schemes of the embedding matrices have been introduced in (Weston et al., 2015):

1. **Adjacent:** the output embedding matrix in the k^{th} hop is shared with the input embedding matrix in the $(k+1)^{\text{th}}$ hop, i.e., $A^{k+1} = C^k$ for $k \in \{1, K-1\}$. Also, the weight matrix W in Equation (4) is shared with the output embedding matrix in the last memory hop such that $W^\top = C^K$.
2. **Layer-wise:** all the weight matrices A^k and C^k are shared across different hops, i.e., $A^1 = A^2 = \dots = A^K$ and $C^1 = C^2 = \dots = C^K$.

In the next section, we show how the task of dialog state tracking can be formalized as machine reading task and solved using such memory enhanced model.

3.3 Dialog Reading Model for State Tracking

In this section, we formalize dialog state tracking using the paradigm of machine reading. As far as our knowledge goes, it is the first attempt to apply this approach and develop a specific dataset format, detailed in Section 4, from an existing and publicly available dialog state tracking challenge dataset to fulfill this task. Assuming (1) a dyadic dialog d composed of a list of utterances, (2) a state composed with (2a) a set of variables v_i with $i = \{1, \dots, n\}$ and (2b) a set of corresponding assigned values l_i . One can define a question q_v that corresponds to the specific querying of a variable in the context of a dialog $p(l_i|q_v, d)$. In such context, a dialog state tracking task consists in determining for each variable v , $l^* = \arg \max_{l_i \in L} p(l_i|q_v, d)$, with L the specific domain of expression of a variable v_i .

In addition to the actual dataset, we propose to investigate four general reasoning tasks using *DSTC-2* dataset as a starting point. In such way, we leverage the dataset of *DSTC-2* to create more complex reasoning task than the ones present in the original dialogs of the dataset by performing rule-based modification over the corpus. Obviously, the goal is to develop resolution algorithms

that are not dedicated to a specific reasoning task but inference models that will be as generic as possible. In the rest of the section, each of the reasoning tasks associated with dialog state tracking are described and the generation protocol is explained with examples.

Factoid Questions : This first task corresponds to the current formulation of dialog state tracking. It consists of questions where a previously given a set of supporting facts, potentially amongst a set of other irrelevant facts, provides the answer. This kind of task was already employed in (Weston et al., 2014) in the context of a virtual world. In that sense, the result obtained to such task are comparable with the state of the art approaches.

Yes/No Questions : This task tests the ability of a model to answer true/false type questions like “*Is the food italian ?*”. The conversion of a dialog to such format is deterministic regarding the fact that the utterances and corresponding true states are known at each utterance of a given dialog.

Indefinite Knowledge : This task tests a more complex natural language construction. It tests if statements can be models in order to describe possibilities rather than certainties, as proposed in (Weston et al., 2014). In our case, the answer will be “maybe” to the question “*Is the price-range required moderate ?*” if the slot hasn’t been mentioned yet throughout the current dialog. In the case of state tracking, it will allow to seamlessly deal with unknown information about the dialog state. Concretely, this set of questions and answers are generated has a super-set of the Yes-No Questions set. First, sub-dialog starting from the first utterance of a given dialog are extracted under the condition that a given slot is not informed in the corresponding annotation. Then, a question-answering question is generated.

Counting and Lists/Sets : This last task tests the capacity of the model to perform simple counting operations, by asking about the number of objects with a certain property, e.g. “*How many area are requested ?*”. Similarly, the ability to produce a set of single word answers in the form of a list, e.g. “*What are the area requested ?*” is investigated. Table 1 give an example of each of the question type presented below on a dialog sample of *DSTC-2* corpus.

Inference procedure: Concretely, the current set of utterances of a dialog will be placed into the memory using sentence based encoding and the




Figure 2: Illustration of the proposed MemN2N based state dialog tracker model with 3 hops.

question will be encoded as the controller state at $t = 1$. The answer will be produced using a softmax operation over the answer vocabulary that is supposed fixed. We consider this hypothesis valid in the case of factoid and list questions because the set of value for a given variable is often considered known. In the cases of Yes/No and Indefinite knowledge question, {Yes, No, Maybe} are added to the output vocabulary. Following (Weston et al., 2014), a list-task answer will be considered as a single element in the answer set and the count question. A possible alternative would be to change the activation function used at the output of the MemN2N from softmax activation function to a logistic one and to use a categorical cross entropy loss. A drawback of such alternative would be the necessity of cross-validating a decision threshold in order to select a eligible answers. Concerning the individual numbers for the count question set, the numbers founded on the training set are added into the vocabulary.

We believe more reasoning capabilities need to be explore in the future, like spacial and temporal reasoning or deduction as suggested in (Weston et al., 2015). However, it will probably need the development of a new dedicated resource. Another alternative could be to develop a question-answering annotation task based on a dialog corpus where reasoning task are present. The closest work to our proposal that can be cited is (Bordes and Weston, 2016). In this paper, the authors defines a so-called End-to-End learnable dialog system to infer an answer from a finite set of eligible answers w.r.t the current list of utterances of the dialog. The authors generate 5 artificial tasks of dialog. However the reasoning capabilities are not explicitly addressed and the author explicitly claim that the resulting dialog system is not satisfactory

yet. Indeed, we believe that having a proper dialog state tracker where a policy is built on top can guarantee dialog achievement by properly optimizing a reward function throughout a explicitly learnt dialog policy. In the case of proper end-to-end systems, the objective function is still not explicitly defined (Serban et al., 2015) and the resulting systems tend to be used in the context of chat-oriented and non-goal oriented dialog systems. In the next section, we present experimental details and results obtained on the basis of the *DSTC-2* dataset and its conversion to the four mentioned reasoning tasks.

4 Experiments

4.1 Dataset and Data Preprocessing

In the *DSTC-2* dialog corpus, a user queries a database of local restaurants by interacting with a dialog system. A dialog proceeds as follows: first, the user specifies constraints concerning the restaurant. Then, the system offers the name of a restaurant that satisfies the constraints. Finally, the user accepts the offer and requests additional information about the accepted restaurant. In this context, the dialog state tracker should be able to track several types of information that compose the state like the geographic area, the food type and the price range slots. In order to make comparable experiments, sub-dialogs generated from the first utterance to each utterance of each dialog of the corpus have been generated. The corresponding question-answer pairs have been generated using the annotated state for each of the sub-dialog. In the case of factoid question, this setting allows for fair comparison at the utterance-level state tracking gains with the prior art. The same protocol has been adopted for the generated reasoning task. In that sense, the tracker task consists

Index	Actor	Utterance
1	Cust	Im looking for a cheap restaurant in the west or east part of town.
2	Agent	Thanh Binh is a nice restaurant in the west of town in the cheap price range.
3	Cust	What is the address and post code.
4	Agent	Thanh Binh is on magdalene street city centre.
5	Cust	Thank you goodbye.
6		Factoid Question What is the pricerange ? Answer: {Cheap}
7		Yes/No Question Is the Pricerange Expensive ? Answer: {No}
8		Indefinite Knowledge Is the FoodType chinese ? Answer: {Maybe}
8		Listing task What are the areas ? Answer: {West,East}

Table 1: : Dialog state tracking question-answering examples from *DSTC2* dataset

in finding the value l^* as defined in Section 3.3. In the overall dialog corpus, Area slot counts 5 possible values, Food slot counts 91 possible values and Pricerange slot counts 3 possible values. In order to exhibit reasoning capability of the proposed model in the context of dialog state tracking, three other dataset have been automatically generated from the dialog corpus in order to support 3 capabilities of reasoning described in Section 3.3. Dialog modification has been required for two reasoning tasks, *List* and *Count*. Two types of rules have been developed to automatically produce modified dialogs. On a first hand, string matching has been performed to determine the position of a slot values in a given utterance and an alternative statement has been produced as a substitution. For example, the utterance “*I’m looking for a chinese restaurant in the north*” can be replaced by “*I’m looking for a chinese restaurant in the north or the west of town*”. A second type of modification has been performed in an inter-utterance fashion. For example, assuming a given value “north” has been informed in the current state of a given dialog, one can add lately in the dialog a remark like “I would also accept a place east side of town”. This kind of statement tends to not affect the overall flow of the dialog and allows to add richer semantic to the dialog. In the future, we plan to develop a richer set of generation procedures to augment the dataset. Nevertheless, we believe this simple dialog augmentation strategy allows to exhibit the competency of the proposed model beyond factoid questions.

4.2 Training Details

As suggested in (Sukhbaatar et al., 2015), 10% of the set was held-out to form a validation set for hyperparameter tuning. Concerning the utterance encoding, we use the so-called *Temporal Encoding* technique. In fact, reading tasks require some notion of temporal context. To enable the model

to address them, the memory vector is modified as such $m_i = \sum_j Ax_{ij} + T_A(i)$, where $T_A(i)$ is the i^{th} row of a dedicated matrix T_A that encodes temporal information. The output embedding is augmented in the same way with a matrix T_c (e.g. $c_i = \sum_j Cx_{ij} + T_C(i)$). Both T_A and T_C are learned during training in an end-to-end fashion. They are also subject to the same sharing constraints as A and C . The embedding matrix A and B are initialized using GoogleNews word2vec embedding model (Mikolov et al., 2013). Also suggested on (Sukhbaatar et al., 2015), utterances are indexed in reverse order, reflecting their relative distance from the question so that x_1 is the last sentence of the dialog. Furthermore, adjacent weight tying schema has been adopted. Learning rate η is initially assigned a value of 0.005 with exponential decay applied every 25 epochs by $\eta/2$ until 100 epochs are reached. Then, *linear start* is used in all our experiments as proposed by (Sukhbaatar et al., 2015). More precisely, the softmax function in each memory layer is removed and re-inserted after 20 epochs. Batch size is set to 16 and gradients with an L_2 norm larger than 40 are divided by a scalar to have norm 40. All weights are initialized randomly from a Gaussian distribution with zero mean and $\sigma = 0.1$. In all our experiments, we have tested a set of the embedding size $d \in \{20, 40, 60\}$. After validation, each model uses a 5-hops depth configuration.

4.3 Experimental results

Table 3 presents tracking accuracy obtained for three variables of the *DSTC2* dataset formulated as *Factoid Question* task. We compare with two established utterance-level discriminative neural trackers, a Recurrent Neural Network (RNN) model (Henderson et al., 2014a) and the Neural Belief Tracker (Mrksic et al., 2016). As suggested in this last work, the first RNN baseline model uses no semantic (i.e. synonym) dictionary, while

Locutor	Utterance	Hop 1	Hop 2	Hop 3	Hop 4	Hop 5
Cust	Im looking for a cheap restaurant that serves chinese food	0.00	0.18	0.11	0.04	0.00
Agent	What part of town do you have in mind	0.33	0.30	0.00	0.00	0.00
Cust	I dont care	0.00	0.00	0.17	0.37	1.00
Agent	Rice house serves chinese food in the cheap price range	0.01	0.00	0.00	0.00	0.00
Cust	What is the address and telephone number	0.58	0.09	0.01	0.00	0.00
Agent	Sure rice house is on mill road city centre	0.03	0.00	0.00	0.00	0.00
Cust	Phone number	0.00	0.00	0.00	0.00	0.00
Agent	The phone number of rice house is 765-239-09	0.02	0.01	0.00	0.00	0.00
Cust	Thank you good bye	0.02	0.42	0.71	0.59	0.00
What is the area ? Answer: dontcare						

Table 2: Attention shifting example for the *Area* slot from *DSTC2* dataset, the values corresponds the p_i values affected to each memory block m_i at each hop of the MemN2N

the improved baseline uses a hand-crafted semantic dictionary designed for the *DSTC2* ontology. In this context, a MemN2N model allows to obtain competitive results with the most close, non-memory enhanced, state of the art approach of recurrent neural network with word embedding as prior knowledge.

Model	Area	Food	Price	Joint
RNN - no dict.	0.92	0.86	0.86	0.69
RNN + sem. dict.	0.91	0.86	0.93	0.73
NBT-DNN	0.90	0.84	0.94	0.72
NBT-CNN	0.90	0.83	0.93	0.72
MemN2N($d = 40$)	0.89	0.88	0.95	0.74

Table 3: **One supporting fact task** : Acc. obtained on DSTC2 test set

As a second result, Table 4 presents the performance obtained for the four reasoning tasks. The obtained results lead us to think that MemN2N are a competitive alternative for the task dialog state tracking but also increase the spectrum of definition of the general dialog state tracking task to machine reading and reasoning. In the future, we believe new reasoning capabilities like spacial and temporal reasoning and deduction should be exploited on the basis of a specifically designed dataset.

5 Conclusion and Further Work

This paper describes a novel method of dialog state tracking based on the paradigm of machine reading and solved using MemN2N, a memory-enhanced neural network architecture. In this context, a dataset format inspired from the current datasets of machine reading tasks has been developed for this task. It is the first attempt to solve this classic sub-problem of dialog management in

Variable	d	Yes-No	I.K.	Count.	List.
Food	20	0.85	0.79	0.89	0.41
	40	0.83	0.84	0.88	0.42
	60	0.82	0.82	0.90	0.39
Area	20	0.86	0.83	0.94	0.79
	40	0.90	0.89	0.96	0.75
	60	0.88	0.90	0.95	0.78
PriceRange	20	0.93	0.86	0.93	0.83
	40	0.92	0.85	0.90	0.80
	60	0.91	0.85	0.91	0.81

Table 4: **Reasoning tasks** : Acc. on DSTC2 reasoning datasets

such way. Beyond the experimental results presented in the experimental section, the proposed approach offers several advantages compared to state of the art methods of tracking. First, the proposed method allows to perform tracking on the basis of segment-dialog-level annotation instead of utterance-level one that is commonly admitted in academic datasets but tedious to produce in a large scale industrial environment. Second, we propose to develop dialog corpus requiring reasoning capabilities to exhibit the potential of the proposed model. In future work, we plan to address more complex tasks like spatial and temporal reasoning, sorting or deduction and experiment with other memory enhanced inference models. Indeed, we plan to experiment and compare the same approach with Stacked-Augmented Recurrent Neural Network (Joulin and Mikolov, 2015) and Neural Turing Machine (Graves et al., 2014) that sounds also promising for these family of reasoning tasks.

References

Jonathan Berant, Vivek Srikumar, Pei-Chun Chen, Abby Vander Linden, Brittany Harding, Brad

- Huang, Peter Clark, and Christopher D. Manning. 2014. Modeling biological processes for reading comprehension. In Alessandro Moschitti, Bo Pang, and Walter Daelemans, editors, *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing, EMNLP 2014, October 25-29, 2014, Doha, Qatar, A meeting of SIGDAT, a Special Interest Group of the ACL*. ACL.
- Antoine Bordes and Jason Weston. 2016. Learning end-to-end goal-oriented dialog. *CoRR*.
- Harry Bunt, Jan Alexandersson, Jean Carletta, Jae-Woong Choe, Alex Chengyu Fang, Koiti Hasida, Kiyong Lee, Volha Petukhova, Andrei Popescu-Belis, Laurent Romary, Claudia Soria, and David Traum. 2010. Towards an ISO standard for dialogue act annotation. In *Proceedings of the Seventh International Conference on Language Resources and Evaluation (LREC'10)*. European Language Resources Association (ELRA), may.
- Oren Etzioni, Michele Banko, and Michael J. Cafarella. 2007. Machine reading. In *AAAI Spring Symposium: Machine Reading*. AAAI.
- Milica Gasic and Steve Young. 2011. Effective handling of dialogue state in the hidden information state POMDP-based dialogue manager. *TSLP*, 7(3):4.
- Zoubin Ghahramani and Michael I. Jordan. 1997. Factorial hidden Markov models. *Machine Learning*, 29(2-3):245–273.
- Alex Graves, Greg Wayne, and Ivo Danihelka. 2014. Neural Turing machines. *CoRR*.
- Matthew Henderson, Blaise Thomson, and Steve Young. 2013. *Proceedings of the SIGDIAL 2013 Conference*, chapter Deep Neural Network Approach for the Dialog State Tracking Challenge, pages 467–471. Association for Computational Linguistics.
- M. Henderson, B. Thomson, and S. J. Young. 2014a. Robust dialog state tracking using delexicalised recurrent neural networks and unsupervised adaptation. In *Proceedings of IEEE Spoken Language Technology*.
- Matthew Henderson, Blaise Thomson, and Jason D. Williams. 2014b. The third dialog state tracking challenge. In *SLT*, pages 324–329. IEEE.
- Matthew Henderson, Blaise Thomson, and Steve Young. 2014c. Word-based dialog state tracking with recurrent neural networks. In *Proceedings of SIGDial*.
- Karl Moritz Hermann, Tomas Kocisky, Edward Grefenstette, Lasse Espeholt, Will Kay, Mustafa Suleyman, and Phil Blunsom. 2015. Teaching machines to read and comprehend. *CoRR*.
- Armand Joulin and Tomas Mikolov. 2015. Inferring algorithmic patterns with stack-augmented recurrent nets. In Corinna Cortes, Neil D. Lawrence, Daniel D. Lee, Masashi Sugiyama, and Roman Garnett, editors, *Advances in Neural Information Processing Systems 28: Annual Conference on Neural Information Processing Systems 2015, December 7-12, 2015, Montreal, Quebec, Canada*, pages 190–198.
- Ankit Kumar, Ozan Irsoy, Jonathan Su, James Bradbury, Robert English, Brian Pierce, Peter Ondruska, Ishaan Gulrajani, and Richard Socher. 2015. Ask me anything: Dynamic memory networks for natural language processing. *CoRR*.
- Angeliki Metallinou, Dan Bohus, and Jason Williams. 2013. Discriminative state tracking for spoken dialog systems. In *Association for Computer Linguistics*, pages 466–475. The Association for Computer Linguistics.
- Tomas Mikolov, Ilya Sutskever, Kai Chen, Gregory S. Corrado, and Jeffrey Dean. 2013. Distributed representations of words and phrases and their compositionality. In Christopher J. C. Burges, Leon Bottou, Zoubin Ghahramani, and Kilian Q. Weinberger, editors, *Advances in Neural Information Processing Systems 26: 27th Annual Conference on Neural Information Processing Systems 2013. Proceedings of a meeting held December 5-8, 2013, Lake Tahoe, Nevada, United States*, pages 3111–3119.
- Diego H. Milone and Antonio J. Rubio. 2003. Prosodic and accentual information for automatic speech recognition. *IEEE Transactions on Speech and Audio Processing*, 11(4):321–333.
- Nikola Mrksic, Diarmuid O Seaghda, Tsung-Hsien Wen, Blaise Thomson, and Steve J. Young. 2016. Neural belief tracker: Data-driven dialogue state tracking. *CoRR*, abs/1606.03777.
- Hoifung Poon and Pedro M. Domingos. 2010. Machine reading: A “killer app” for statistical relational AI. In *Statistical Relational Artificial Intelligence*, volume WS-10-06 of *AAAI Workshops*. AAAI.
- Antoine Raux and Yi Ma. 2011. Efficient probabilistic tracking of user goal and dialog history for spoken dialog systems. In *INTERSPEECH*, pages 801–804. ISCA.
- Matthew Richardson, Christopher J. C. Burges, and Erin Renshaw. 2013. MCTest: A challenge dataset for the open-domain machine comprehension of text. In *EMNLP*, pages 193–203. ACL.
- Iulian Vlad Serban, Ryan Lowe, Laurent Charlin, and Joelle Pineau. 2015. A survey of available corpora for building data-driven dialogue systems. *CoRR*.
- Sainbayar Sukhbaatar, Arthur Szlam, Jason Weston, and Rob Fergus. 2015. End-to-end memory networks. In Corinna Cortes, Neil D. Lawrence,

- Daniel D. Lee, Masashi Sugiyama, and Roman Garnett, editors, *Advances in Neural Information Processing Systems 28: Annual Conference on Neural Information Processing Systems 2015, December 7-12, 2015, Montreal, Quebec, Canada*, pages 2440–2448.
- Blaise Thomson and Steve Young. 2010. Bayesian update of dialogue state: A POMDP framework for spoken dialogue systems. *Computer Speech & Language*, 24(4):562–588.
- Laurens van der Maaten and Geoffrey Hinton. 2008. Visualizing data using t-SNE. *Journal of Machine Learning Research*, 9:2579–2605, November.
- Tsung-Hsien Wen, Milica Gasic, Nikola Mrksic, Lina Maria Rojas-Barahona, Pei-Hao Su, Stefan Ultes, David Vandyke, and Steve J. Young. 2016. A network-based end-to-end trainable task-oriented dialogue system. *CoRR*.
- Jason Weston, Sumit Chopra, and Antoine Bordes. 2014. Memory networks. *CoRR*.
- Jason Weston, Antoine Bordes, Sumit Chopra, and Tomas Mikolov. 2015. Towards AI-complete question answering: A set of prerequisite toy tasks. *CoRR*.
- Jason D. Williams, Pascal Poupart, and Steve Young. 2005. Factored partially observable markov decision processes for dialogue management. In *In 4th Workshop on Knowledge and Reasoning in Practical Dialog Systems*, pages 76–82.
- Jason D. Williams, Antoine Raux, and Matthew Henderson. 2016. The dialog state tracking challenge series: A review. *D&D*, 7(3):4–33.
- Jason D. Williams. 2014. Web-style ranking and slu combination for dialog state tracking. In *Proceedings of SIGDIAL*. ACL Association for Computational Linguistics, June.
- Peter Z. Yeh, Benjamin Douglas, William Jarrold, Adwait Ratnaparkhi, Deepak Ramachandran, Peter F. Patel-Schneider, Stephen Laverty, Nirvana Tikku, Sean Brown, and Jeremy Mendel. 2014. A speech-driven second screen application for TV program discovery. In Carla E. Brodley and Peter Stone, editors, *Proceedings of the Twenty-Eighth AAAI Conference on Artificial Intelligence, July 27 -31, 2014, Québec City, Québec, Canada*, pages 3010–3016. AAAI Press.